

A novel method for pulmonary emboli visualization from high-resolution CT images

Eric Pichon^{a*}, Carol L. Novak^b, Atilla P. Kiraly^b, David P. Naidich^c

^aSchool of Electrical & Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332

^bSiemens Corporate Research, 755 College Road East, Princeton, NJ 08540

^cRadiology Dept., New York University Medical Center, 560 First Avenue, New York, NY 10016

ABSTRACT

Pulmonary Embolism (PE) is one of the most common causes of unexpected death in the US. The recent introduction of 16-slice Computed Tomography (CT) machines allows the acquisition of very high-resolution datasets. This has made CT a more attractive means for diagnosing PE, especially for previously difficult to identify small subsegmental peripheral emboli. However, the large size of these datasets makes it desirable to have an automated method to help radiologists focus directly on potential candidates that might otherwise be overlooked. We propose a novel method to highlight potential PEs on a 3D representation of the pulmonary arterial tree. First lung vessels are segmented using mathematical morphology techniques. The density values inside the vessels are then used to color the outside of a Shaded Surface Display (SSD) of the vessel tree. As PEs are clots of significantly lower Hounsfield Unit (HU) values than surrounding contrast-enhanced blood, they appear as salient contrasted patches in this 3D rendering. During preliminary testing on 6 datasets 19 PEs out of 22 were detected (sensitivity 86%) with 2 false positives for every true positive (Positive Predictive Value 33%).

Keywords: Pulmonary Embolism, Computed Tomography, detection, visualization

1. INTRODUCTION

Pulmonary Embolism (PE) occurs when a blood clot obstructs a pulmonary artery. This life threatening pathology is common, with an annual incidence rate of between 23 and 69 per 100,000. It has been estimated that PE causes the death of up to 60,000 people per year in the US alone, making it one of the most common preventable causes of death in hospitals. The standard treatment for PE is anticoagulant drugs to dissolve the clots.

PE is a difficult disease to diagnose clinically since many other conditions share the same symptoms.¹ Many patients with PE are never studied and the majority of patients suspected of having PE do not have the disease.²⁻³ The prompt and accurate diagnosis is of major concern because untreated PE is potentially fatal, while overdiagnosis may lead to unnecessary interventions, including unnecessary anticoagulant treatment with its independent risks of morbidity and mortality.

Although ventilation-perfusion scintigraphy is frequently employed, it remains an indirect method for assessing the presence of clot. Pulmonary angiography has long been considered the “gold standard” for diagnosing PE. Despite its relative safety, this procedure is still associated with definable morbidity and expense. In addition, considerable inter-observer variability has been documented, especially in the diagnosis of small, subsegmental pulmonary emboli – those of greatest concern when interpreting CT angiograms (CTA).⁴

Unlike pulmonary angiography, CTA is easily performed, rapid and non-invasive. As intravenous contrast material does not penetrate clot, CT allows the direct visualization of clot obstructing a vessel or blood streaming around a partially occluding embolus. Recent studies have shown that CTA has better sensitivity, selectivity and inter-observer agreement than other diagnostic methods, including ventilation perfusion scintigraphy as well as routine pulmonary angiography.⁴⁻⁸ In particular, in a recent animal study with absolute ground truth validating the presence of clot, it has been shown that

* eric@ece.gatech.edu

the angiography “gold standard” yields both false positive and false negative interpretations, with a sensitivity and positive predictive value comparable to CT.⁹

The introduction of 16-slice CT machines, allows the acquisition of near isotropic datasets, typically 0.5 mm in all directions. This has made CT a still more attractive means for diagnosing PE. However these datasets are very large – typically on the order of 512^3 voxels – which puts a heavy burden on radiologists. To this end, there has been some preliminary work in developing algorithms for automatic detection of pulmonary emboli from CT.¹⁰⁻¹¹

In this paper we propose a novel method to highlight potential PEs on a 3D representation of the pulmonary arterial tree. Section 2 reviews some background on signed distance maps. We show in section 3 how these maps will be used to implement morphological operations during the segmentation process, and also how they are used to determine correspondences between points on the outer surface of the arterial tree and points inside the arteries. These correspondences will be used in the proposed scheme to color the surface of a 3D rendering of the arterial tree, thus revealing inside structures such as PEs. In section 4 we show preliminary results on synthetic and real datasets.

2. SIGNED DISTANCE MAPS

In this section we briefly review the theory of Signed Distance Maps (SDM) and their relation with standard mathematical morphology primitives.¹²

Consider an object S in 3D space* and its boundary ∂S . The SDM corresponding to S is the map:

$$D_S : R^3 \rightarrow R$$

$$p \mapsto D_S(p) := \varepsilon_S(p) \min_{q \in \partial S} d(p, q)$$

where $\varepsilon_S(p) = -1$ for $p \in S$ and $\varepsilon_S(p) = 1$ for $p \notin S$ and $d(p, q)$ is the distance between points p and q . For any point $p=(x,y,z)$ in the space, the SDM associates the distance to the closest point on the boundary ∂S of S with a negative sign for points inside S . See figure 1 for an example SDM.



Fig 1a. Initial object S




Fig 1b. Distance map. Darker pixels correspond to higher values of D_S .

Efficient computation of distance transforms is still an active research topic and different approaches have been proposed.¹³ In this work we will use the fact that D_S is the solution of the Eikonal equation:

$$\|\nabla D_S\| = 1 \text{ with boundary condition } D_S(p)|_{p \in \partial S} = 0$$

which can be efficiently solved using the Fast Marching algorithm initially developed for curve and surface evolution. The algorithm is initialized by setting D_S to 0 on the boundary ∂S , and marching away in the normal direction with unit speed.¹⁴⁻¹⁵

* This theory and our approach extend to any dimension, although for this work we concentrate on 3D. Some figures are in 2D for clarity.

A very interesting property is that the level sets of D_S correspond to the erosions and dilations of S with a Euclidean ball of corresponding radius. For example, a threshold of $+2.0$ will correspond to a dilation with a ball of radius 2.0 . A threshold of -1.0 will correspond to an erosion with a ball of radius 1.0 . Figure 2 shows examples.




Fig 2a. Distance map. The boundary ∂S of S is shown in gray.




Fig 2b. Thresholding with a positive value corresponds to a dilation.



Fig 2c. Thresholding with a negative value corresponds to an erosion.

Once the SDM is computed, the cost of a morphological operation does not depend on the size of the structure element. This is particularly useful if we want to use large spheres. This subject is covered in greater detail elsewhere.¹⁶⁻¹⁷

Another interesting property¹⁸ is that the shocks of D_S (the points for which the Eikonal equation is not respected) $M_S = \{p : \|\nabla D_S(p)\| \neq 1\}$ form a medial axis of S as shown in Fig 3a.




Fig 3a. Distance map. The boundary ∂S of S is shown in gray and a medial axis M_S in black. Notice that the exterior skeleton is also visible.




Fig 3b. Distance map and a few paths (represented in black) from the boundary (represented in gray) to the previously defined medial axis

Therefore if we pick a point p on S we can determine a path to the defined medial axis by following the steepest ascent ∇D_S as long as $\|\nabla D_S(p)\| = 1$. Equivalently if we start from a point b on ∂S this corresponds to maximizing D_S while $D_S(p) \approx \|p - b\|$. This is shown in Fig 3b. This defines a correspondence between points on the boundary ∂S and points of S . A simple interpretation is that every point of S is associated with the nearest boundary point on ∂S . Conversely, every point on ∂S is associated with a set of points in S , namely those that fall on the shortest path from the medial axis M_S . The distance criteria ensures that the ascent of D_S will not take place along the medial surface (which could be an ascending path on D_S , an example is the axis of a cone in 3D). This is equivalent to forcing the path to be a straight line. Section 3.2 contains further details on this algorithm.

3. ALGORITHM

In our algorithm, we begin by segmenting lung vessels using mathematical morphology techniques. The inner gray-level values of the vessels are then used to color the outside of the vessel tree. More precisely the color of every surface element will reflect the values of interior voxels that form a path to the corresponding interior-most part of the vessel as proposed in the previous section. Because PEs are clots of significantly lower Hounsfield Unit (HU) values than surrounding contrast-enhanced blood, they appear as salient contrasted patches in this 3D rendering.

3.1. Segmentation

The segmentation proceeds in two steps. We first obtain the segmentation of the lungs and trachea. Next, we use this segmentation as a mask for the segmentation of the arteries within the lungs. In the remainder of this paper, we refer to the original volume as V , the segmented volume of the lungs as M , and the segmented volume of the arteries as S . Additionally, given the voxel location \mathbf{x} , its value within the volume V is referred to as $V(\mathbf{x})$. All voxels within the segmented volumes either have binary values of 0 for outside of the segmentation and 1 for inside of the segmentation.

3.1.1. Coarse segmentation of the lungs

In this section we describe a simple method based on mathematical morphology to create a segmentation of the lung region M , encompassing the pulmonary arteries and veins. Since this segmentation is used as a mask for the arterial segmentation, we refer to it as the lung mask.

We start by segmenting air inside the lungs using region growing: starting from one seed (for example in the trachea) all contiguous voxels of intensity lower than a fixed threshold are added to the segmented volume. The threshold is set high enough to penetrate the airway walls, enter the parenchyma, and fill the lung. Note that this threshold does not cause the segmentation to reach outside of the lungs since the chest wall is usually well defined. Figure 4 shows the segmentation.



Fig 4a. Initial CT image (coronal slice).




Fig 4b. Segmentation of air inside the lungs. The mask is shown in medium gray and overlaid on the original image.

We then perform a dilation and an erosion via Euclidean spheres, as shown in Figure 5. We use a slightly larger sphere for the erosion (12mm) than for the dilation (10mm). The effect is that the resulting mask is slightly inside the lungs. This will ensure that we will not later accidentally classify part of a rib as an artery. In addition, the dilation fills in cavities of the lung mask. These cavities are mainly due to arteries and veins that must be included in this mask.

All of these morphological operations are performed fully in 3D and take into account any anisotropy of the dataset. They are implemented using distance maps as described in the previous section.

Even though this mask is generated mainly as a pre-processing step towards the segmentation of the arterial tree, we also use it during rendering to give the physician anatomical context as demonstrated in Figure 6.




Fig 5a. Result of dilation.




Fig 5b. Result of erosion.




Fig 6a. Segmented lungs




Fig 6b. Segmented lungs and axial and coronal CT slices

3.1.2. Segmentation of the pulmonary arterial tree

This section describes the segmentation of the arterial tree base to produce the volume S . Using the previously generated lung mask M , we threshold the dataset for everything inside the lungs that is not air, and discard the smallest connected components. Specifically, given the minimum threshold T at which the arteries are visible, any voxel \mathbf{x} in V such that $V(\mathbf{x}) > T$ and $M(\mathbf{x}) = 1$ will be included in the segmentation S . All other voxels \mathbf{x} will be excluded, i.e. $S(\mathbf{x}) = 0$. Currently a limitation is that veins are segmented along with arteries. The results are shown in Figure 7.

3.2. Surface rendering of volumetric data

The values *inside* the vessels are then used to color the *outside* of the 3D Shaded Surface Display (SSD) of the vessel tree. When the 3D display is rendered, the color of every surface element is calculated from the voxels along the path from that surface element towards the medial axis of the artery. Since PEs appear as clots of significantly lower density values than surrounding contrast-enhanced blood, paths that pass through PEs will have lower values than those passing through unblocked arteries. In turn the 3D rendering of the entire arterial tree will show clots as patches of obviously different density. This allows for a novel method to map 3D data onto a 2D surface, hence combining the advantages of surface rendering and 3D data analysis.




Fig 7a. Segmented tree (coronal slice).



Fig 7b. Segmented lungs and vascular tree.

The following is the algorithm used to compute the path to the medial axis from a surface voxel \mathbf{b} of the vessel segmentation S . This path is simply a list of voxels from the surface to the medial axis and is determined by the steepest ascent to the central most portion of the vessel. The conditional in Step 3 prevents the path from continuing along the medial axis.

For all voxels \mathbf{b} of ∂S

- 1. $\mathbf{p} := \mathbf{b}$, $\mathbf{C}(\mathbf{b}) = \emptyset$**
- 2. $\mathbf{C}(\mathbf{b}) = \mathbf{C}(\mathbf{b}) \cup \mathbf{V}(\mathbf{p})$**
- 3. if $D_S(p) \approx \|p - b\|$ then**
 - a. move \mathbf{p} by one voxel to maximize D_S**
 - b. jump to 2.**
- 4. We have reached the medial axis.**
Color \mathbf{b} using the values in set $\mathbf{C}(\mathbf{b})$
(e.g. mean value)

4. RESULTS

We have evaluated our method on both synthetic data and real patient data to demonstrate that the method is feasible and useful.

4.1. Synthetic dataset

To illustrate the concept, we constructed a simple 3D digital simulation of branching vessels. HU values are 200 for “contrasted blood” and 50 for 2 spherical “Pulmonary Emboli” that were inserted into the data. We then added noise (uniformly distributed in $[-100, 100]$). The resulting simulation is shown in Fig 8a.




Fig 8a. Synthetic 3D phantom with 2 PEs. (1 slice)

Along each path towards the defined medial axis (see Fig 3b) there will be a range of density values. There are several ways to summarize these values into one. The arithmetic mean is used in Figure 9a. In the case of PE detection we are more interested in lower values, and therefore the 1st quartile on the path might also be used as shown in Figure 9b. The minimum value can also be employed as shown in Figure 9c.

In Figure 9a-c, the SSDs visualizing the outside of the artery readily show both clots inside the synthetic artery. In the original image the clots show as dark patches within brighter areas of contrast-filled blood. However for the purposes of gray scale rendering, the visualization is inverted, with lower CT values rendered as brighter patches.




Fig 9a. Mean HU values



Fig 9b. 1st quartile of HU values



Fig 9c. Minimum of HU values

Because PEs manifest as clots of significantly lower HU values, one could argue that coloring using the minimum value will yield the best representation. In that case it is enough that the path from the surface of the arterial tree to the medial axis cross only *one* voxel of a PE for the color of the corresponding surface element to be set to that dark value and produce a salient patch. This is illustrated in figure 9c where the 2 artificial PEs are very clearly visible. However this representation is very sensitive to noise. Another possibility is to use the mean value on the path as shown in Fig 9a. This produces more moderate results: patches are less visible, but the representation is much more robust to noise. However in real CT datasets, the contrast agent is not as uniformly distributed as in our synthetic example, and areas of exceptionally high contrast may appear adjacent to clots. In these cases averaging smoothes away the useful information.

Using the 1st quartile of the values on the path to the medial axis is a good tradeoff between these two extremes. At least 25% of the voxels on the path will have to be significantly darker for the visualization to be affected, and high HU values are not taken into account. This addresses both the problem of isolated spurious dark voxels due to noise, and that of regions of high concentrations of the contrast agent in the blood flow.

4.2. Real datasets

In our preliminary experiments we used multi-slice CT datasets of 6 patients referred for possible PE. Three of the patients were negative (no pulmonary emboli at all) and 3 of them were positive. The positive patients had multiple PEs. An experienced radiologist marked PEs in the data using conventional soft-copy reviewing of 2D slices. Altogether 22 PEs in 3 patients were identified. Using a different label, zones that looked suspicious but which did not correspond to PEs (false positives) were also annotated. Most of the physician-marked false positives were in veins or due to turbulent flow effects.

Figure 10a shows a sagittal slice from a positive patient, with 3 PEs indicated by arrows. Figures 10b and 10c show a rendering of the arterial tree from the CT dataset shown in 10a. In 10b the mean value along the path to the arterial axis is used to color the corresponding surface patch. In 10c the 1st quartile of values is used for the rendering.

The datasets were screened using the 1st quartile representation (as shown in Fig 10c) due to its balance of sensitivity and noise. A PE candidate was counted for every large (several dozen voxels) patch of low HU values. We did not take into account small (less than a dozen voxels) isolated bright patches that corresponded only to partial volume effects in small subsegmental vessels. Fig 10d shows the results, where PE candidates are circled.




Fig 10a. CT dataset (sagittal slice). Three PEs are indicated by arrows.




Fig 10b. Proposed visualization using mean value)




Fig 10c. Proposed visualization (1st quartile)




Fig 10d. Resulting candidates

Altogether, 19 out of 22, or 86% of the PEs could be very clearly distinguished. There were on average 2 false positives for every true PE, giving a positive predictive value of 33%. All the zones labeled by the radiologist as suspicious were also detected by our method. They corresponded to roughly half of the computer's false positives.

5. CONCLUSION

We presented a novel method for visualization of pulmonary embolism from high resolution computed tomography datasets. The traditional method of reviewing 3D data by examining adjacent 2D slices requires readers to simultaneously keep track of multiple arterial branches, examining each of them for clots. As ever smaller arteries can be visualized due to increasing data resolution, this task becomes more challenging for radiologists. The technique we present here condenses the volumetric information onto a surface map of the arterial tree, thus reducing the dimension of the visualization problem in an anatomically relevant way. Furthermore, the method is not limited to the detection of pulmonary embolism and could certainly be used to visualize the interior of intricate 3D shapes in other applications. Preliminary clinical results are encouraging and a larger validation study is underway.

REFERENCES

1. Manganelli D, Palla A, Donnataria V, Giuntini C, "Clinical features of pulmonary embolism: doubts and certainties", *Chest*; **107 (1 suppl)**: 22S-32S, 1995.
2. Gosselin MV, Rubin GD, Leung AN, Huang J, Rizk NW, "Unsuspected Pulmonary Embolism: Prospective Detection on Routine Helical CT", *Radiology* **208**: 209-215, 1998.
3. Winston CB, Wechsler RJ, Salazar AM, Kurtz AB, Spirn PW, "Incidental Pulmonary Emboli Detected at Helical CT: Effect on Patient Care", *Radiology* **201**: 23-27, 1996.
4. "Value of the ventilation/perfusion scans in acute pulmonary embolism: results of the Prospective Investigation of Pulmonary Embolism Diagnosis (PIOPED) the PIOPED Investigators", *JAMA* **263**: 2753-2759, 1990.
5. Mayo JR, Remy-Jardin M, Muller NL, Remy J, Worsley DF, Hossein-Foucher C, Kwong JS, Brown MJ, "Pulmonary Embolism: Prospective Comparison of Spiral CT with Ventilation-Perfusion Scintigraphy", *Radiology*; **205**: 447-452, 1997.
6. Remy-Jardin M, Remy J, Deschildre F, Artaud D, Beregi JP, Hossein-Foucher C, Marchandise X, Duhamel A, "Diagnosis of Pulmonary Embolism with Spiral CT: Comparison with Pulmonary Angiography and Scintigraphy", *Radiology* **200**:699-706, 1996.
7. Chartrand-Lefebvre C, Howarth N, Lucidarme O, Beigelman C, Cluzel P, Mourey-Gerosa I, Cadi M, Grenier P, "Contrast-enhanced helical CT for pulmonary embolism detection: inter- and intraobserver agreement among radiologists with variable experience", *AJR* **172**:107-112, 1999.
8. van Rossum AB, Pattynama PM, Ton ER, Treurniet FE, Arndt JW, van Eck B, Kieft GJ, "Pulmonary Embolism: Validation of Spiral CT Angiography in 149 Patients", *Radiology* **201**:467-470, 1996.
9. Baile EM, King GG, Muller NL, D'Yachkova Y, Coche EE, Pare PD, Mayo JR, "Spiral computed tomography is comparable to angiography for the diagnosis of pulmonary embolism", *Am J Respir Crit Care Med* **161(3 Pt 1)**: 1010-1015, 2000.
10. Masutani Y, MacMahon H, Doi K, "Computerized detection of pulmonary embolism in spiral CT angiography based on volumetric image analysis", *IEEE Transactions on Medical Imaging* **21(12)**: 1517-1523, 2002.
11. Zhou C, Hadjiiski LM, Sahiner B, Chan HP, Patel S, Cascade PN, Kazerooni EA, Wei J, "Computerized detection of pulmonary embolism in 3D computed tomographic images: vessel tracking and segmentation techniques", *Medical Imaging 2003: Image Processing*, 5032:1613-1620, SPIE, San Diego, 2003.
12. Gonzalez R, Woods RE, *Digital Image Processing*, Addison-Wesley, Reading, MA, 2002
13. Cuisenaire O, "Distance transformations: fast algorithms and applications to medical image processing", PhD thesis, Universite Catholique de Louvain, Belgium, 1999.
14. Sethian JA, *Level Set Methods and Fast Marching Methods*, Cambridge University Press, Cambridge, UK, 1999.
15. Breen DE, Mauch S, Whitaker RT "3D Scan Conversion of CSG Models into Distance, Closest-Point and Colour Volumes", *Volume Graphics*, Springer, London, 135-158, 2000.
16. Arhart AB, Vincent L, Kimia BB, "Mathematical Morphology: The Hamilton-Jacobi Connection" *Proceedings of the International Conference on Computer Vision*, 215-219, IEEE Computer Society, Berlin, 1993.
17. Sapiro G, Kimmel, Shaked D, Kimia BB, Bruckstein AM, "Implementing Continuous-Scale Morphology Via Curve Evolution", *Pattern Recognition*, **26**: 1363-1372, 1993.
18. Leymarie F, Levine MD, "Simulating the Grassfire Transform Using an Active Contour Model", *PAMI* (**14**) **1**: 56-75, 1992.